A Case Study on Visual Analytics for Optimizing Drug Duplicate Alerts in a Medication Clinical Decision Support System

Jaehoon Lee, PhD[1,2], Wendi L. Record, PharmD[1], Nathan C. Hulse, PhD[1,2]
[1]Intermountain Healthcare, Salt Lake City, UT; [2]University of Utah, Salt Lake City, UT

Abstract

In this study, we proposed and implemented a visual analytics approach to reduce drug duplicate alerts in a medication clinical decision support system running in our local hospitals. Our approach includes 1) development of a dashboard to provide interactive visualization of time-driven alert data, 2) statistical analysis of effectiveness of duplicate alerts based on behavioral change of physicians by alerts, and 3) process monitoring that enables detection of abnormal patterns in alert generation over time. Based on the tools, we have reduced drug duplicate alerts by removing clinically insignificant alerts (nuisance alert), eliminating duplication between different positions and applications, and detecting broken functions in the decision support system. The results from proof-of-concept implementation shows that the proposed approach could reduce duplicate alerts with increase of alert effectiveness measures: overridden rate and provider behavior change rate.

Keyword

Medication clinical decision support system; Electronic health record; Computerized provider order entry

Introduction

Medication related Clinical Decision Support (mCDS) systems are one of the key components in modern electronic health record (EHR) systems that are specialized in preventing and reducing human errors related to drug prescription [1]. They have been widely adopted in the clinical fields of healthcare institutes and are known to have a positive impact on preventing adverse drug events [2-4]. Integrated with computerized physician order entry (CPOE), mCDS supports physicians ordering medications with minimized human errors with functions of drug-allergy checking, basic dosing guidance, formulary decision support, duplicate therapy checking, and drug–drug interaction checking, etc. These decision supports may be delivered to providers as an intervention to recommend change or reconsider of their action, typically as a form of “alert” in their computer systems.

Despite these benefits, mCDS is also known of its limitation that it often generates too many alerts when it is configured by rigorous rules, and results in physical and mental burdens to providers. When an alert is presented in provider’s desktop, it is asked to recognize it and take actions to override it or follow suggested actions by the system, optionally or mandatorily. It is known that too many alerts result in apathy of providers against alerts: alert fatigue, that is a status where alerts are ignored regardless of their effectiveness. Therefore, the goal of mCDS alert optimization is to minimize the number of alerts presented to users while maintaining or maximizing the effectiveness [5,6].

The basic idea of optimizing alerts is identifying and removing relatively ineffective alerts with expectation of reduced total number of alerts and increased total effectiveness. Although a number of studies have tried to optimize mCDS alerts in a systematic way, a limitation often lies in accurately measuring effectiveness. The effectiveness of mCDS alerts can be affected by various factors such as clinical contexts on how and why alerts are generated, clinical settings, whether an alert is accepted or overridden, and characteristics of providers seen by [7,8]. Because of the features, optimizing mCDS alerts in a data driven approach used to be a manual and time-consuming work accompanying repetitive data extraction and analysis and discussion between interdisciplinary domain experts.

With recent advances in the area of data science technologies, a visual analytics approach supporting agile optimization using automated data extraction and interactive visualization tools was presented [9]. By extending the idea, we introduce our approach of visual analytics that integrates several efforts of optimizing mCDS alerts, by 1) an interactive dashboard of time-driven alert data, 2) statistical analysis of alert effectiveness, and 3) statistical process monitoring that enables detection of abnormal patterns in alerts over time. We selected drug duplicate alerts in our study to validate the feasibility of this approach, since they have been investigated less than other major alert types (e.g. drug-drug interaction) and are known to be hard to be optimized due to their complicated nature. For example, drug duplicate alerts can be reduced by eliminating redundant alerts, which implies that a same alert is generated and seen unnecessarily multiple times by different providers or in different clinical applications. We believe our visual analytics approach enables to analyze data with different aspects in clinical workflows so that we can track such redundancy in an intuitive way.

We developed a dashboard employs the proposed functions and runs based on the mCDS system in our local hospitals. We built a team to review data from the dashboard regularly and have found several cases causing excessive alerts. We have repeated the cycle of improvement for six months and pursued actions to eliminated the causes. As a result, it shows that the proposed approach successfully reduced drug duplicate alerts while maintaining key effectiveness measures: alert overridden rate and provider behavior change rate.

Background

Intermountain Healthcare is a not-for-profit integrated delivery network that serves the populations in the area of the Intermountain West (Utah and southern Idaho). It has 22 hospitals, over 150 clinics, a medical group of over 700 employed physicians and an insurance plan that serves the needs of the people in the region. Recently, Intermountain has implemented a new EHR system: iCentra, in partnership with the Cerner Corporation. A key part of the new system is mCDS, which is designed to help providers with decision support functions to prevent medication related orders. The mCDS consists of a rule engine to store and translate computer interpretable logics and an event handler that fires an alert based on input data and the logics. The mCDS is integrated with CPOE so that such decision support functions are integrated with clinical workflow of medication ordering processes. Most of mCDS rules are executed whenever a new medication order (triggering order) is entered in the CPOE.

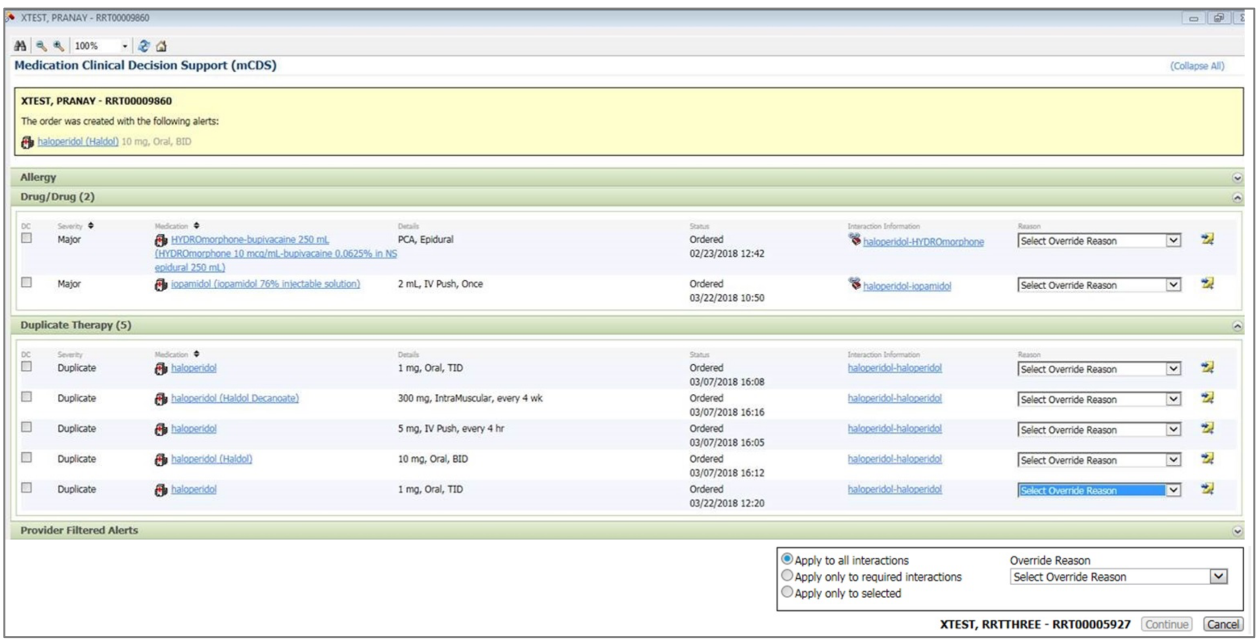

Figure 1. Screenshot: mCDS in iCentra

mCDS alerts are delivered through an alert dialog inside a provider's desktop (See Figure 1). Four alert types are supported by mCDS: allergy-drug, drug-allergy, drug-drug interaction, and drug duplicate. Since one triggering order can be associated with multiple orders already made for a patient (i.e. precondition order) at the time of ordering, an alert dialog may consist of multiple alert sections for each represents association between a triggering order and precondition orders. In addition to that, mCDS also allows "suppression" that is a function to block alerts depending on specific conditions. Once an alert dialog shows up, a provider can choose to continue or remove a triggering alert (See Figure 1 bottom right). Also the provider can discontinue precondition orders by clicking check boxes on the right in the dialog and optionally enter overridden reason for each precondition order by selecting from the list or manually entering free text.

Drug duplicate alerts are designed to detect inappropriate duplication of therapeutic groups or active ingredients and are estimated significant proportion of volumes in medication related alerts [10]. In our mCDS, we use Multum drug category as the therapeutic group, which is a drug ontology and database developed by Cerner. To produce a duplicate alert, mCDS checks if a triggering order fall under the same medication category with already placed orders in the patient's encounter. However it is hard to optimize duplicate alerts, as their nature is related to clinical workflow or logistics processes, such as outpatients receiving prescriptions from different prescribers or early refill sue to holidays [11]. We selected drug duplicate alerts to validate the usefulness of the proposed approach, as visual analytics provide a way to analyze data with different aspects in complicated clinical workflows so that we can detect inappropriate alerting pattern in an intuitive way.

Method

Key measures: We built a team consists of medical informaticists, pharmacy informaticists, data analysts, business intelligence developers to derive key measures of mCDS duplicate alerts and develop design concept of the dashboard. Based on analysis from the team about the front-end workflow in Figure 1, we derived key measures to monitor as Table 1. In addition, we added contextual information such as facility, care / nursing unit, provider position, etc. and different aggregation levels by time, such as weekly averaged measures from Table 1. We also developed normalized measures to incorporate the volume of medication orders.

Table 1. Key measures related to duplicate alert

Level	Measure
Alert dialog	# of alert dialog seen by user # of alert dialog with continued triggering order # of alert dialog with removed triggering order # of alert dialog with modification of at least one precondition orders within 10 minutes
Precondition orders	# of alert fired / generated in an alert dialog # of alert overridden reason entered (either selected or typed) # of alert suppressed by system # of modification of precondition orders

Effectiveness measure: We developed two outcome measures for effectiveness of alerts: overridden rate and behavioral rate. An alert is considered to be overridden when the alert is recognized by providers but doesn't result in any change of orders. In our mCDS, overridden rate can be measured by counting the number of overridden reason entered as the system allows users enter overridden reason in an alert dialog. Realistically not all overridden alerts have been entered of reasons, therefore "% overridden reason entered" may be an underestimated measure. We developed another measure: behavioral change rate that is defined by frequency of modification of triggering and/or precondition orders by an alert. Such modification can occur while an alert dialog is presented or afterward. Based on the literatures we consider modification within 10 minutes after an alert as a behavioral change by the alert [12,13]. Below are definitions of the two measures:

$$\%\ \text{Behavioral change} =$$

$$\frac{\#\ \text{of alert dialog with triggering order removed} + \#\ \text{of alert dialog with precondition order modified within 10 mins}}{\#\ \text{of total alert dialog}}$$

$$\%\ \text{Overridden reason entered} = \frac{\#\ \text{of alert with overriden reason entered}}{\#\ \text{of total alert dialog}}$$

Dashboard development: We used Enterprise Data Warehouse (EDW) at Intermountain Healthcare as a data source, which is an integrated clinical data source that maintains copies of data from clinical applications in our EHR system for the research and quality improvement purposes [14]. We developed structured query language (SQL) based queries to extract data from the EDW and calculate the measures above. Using the queries, we developed a Tableau dashboard that runs on the web within internal network at Intermountain (See Figure 2).

The dashboard consists of a context filter and three charts; 1) a line chart with daily volume of alert dialog, 2) a line chart with daily volume of rule fired (the number of alerts in an alert dialog), 3) and a bar chart with daily volume of medication orders. On the right of the dashboard, a context filter is provided for users to narrow down the data by encounter types, alerted providers, provider positions, units and facilities, drug category of triggering and precondition orders, order set association, overridden reason, and suppression types. A filter to select measure types (volume / normalized volume) are also provided.

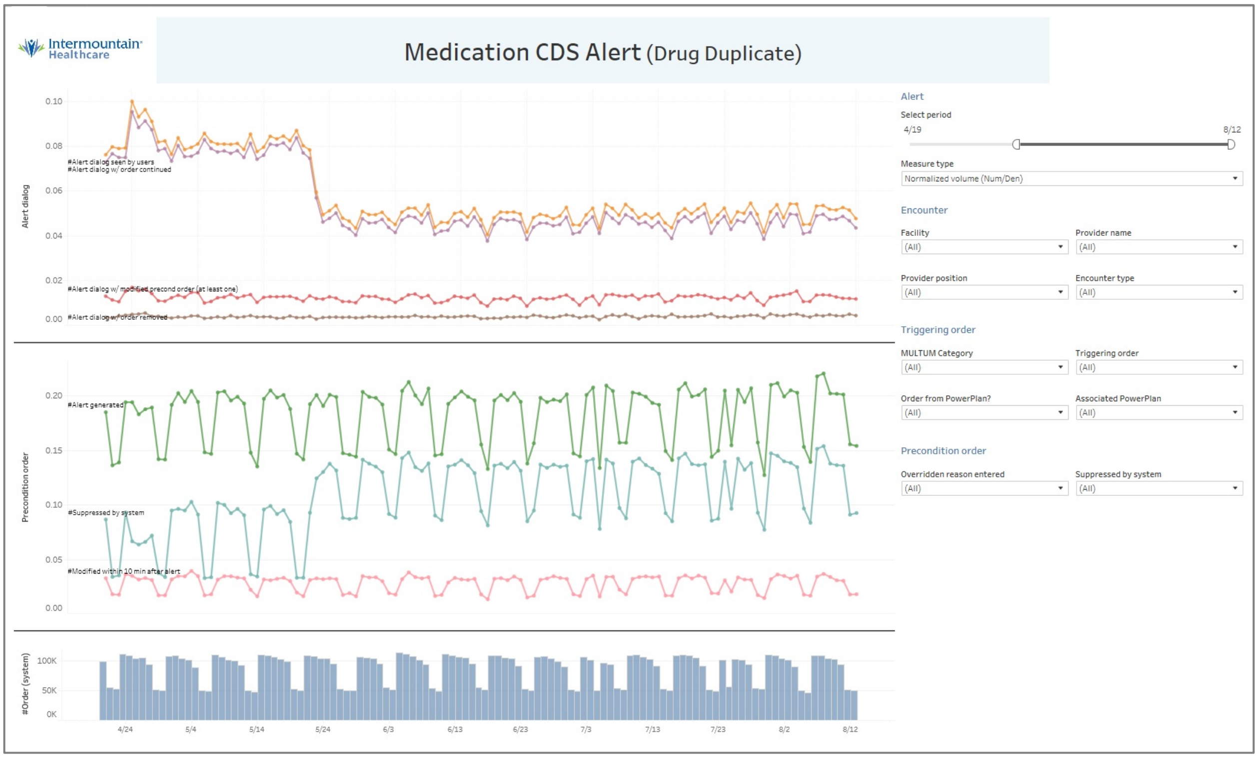

Figure 2. Screenshot: main view of the mCDS dashboard

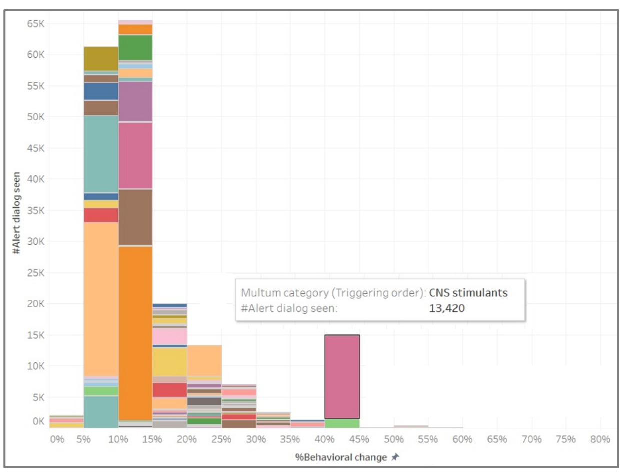

Figure 3. Histogram sorted by effectiveness of duplicate alerts (% of behavioral change)

Figure 4. An email message created by Tableau server to notify abnormal pattern

Figure 3 is another view in the dashboard with a histogram of alert volume. Color represents drug categories of triggering orders and size of bars represents volume of alerts in the categories. X axis represents bins of behavioral change rate and Y axis represents volume of alert dialog showed. In Figure 3, selected bar (purple) and its tooltip show that duplicate alerts in the CNS stimulants category were fired 13,420 times and changed user behaviors by 40-45%. This histogram may be useful to analyze effectiveness of duplicate alerts, by 1) which drug categories were fired of alerts frequently and 2) how much they changed provider behaviors.

In addition to the views, we set up a monitoring function in the Tableau server, that runs if a daily volume of alerts is smaller or bigger than the same weekday of previous week to certain level (e.g. ±10%) and send a notifying email to assigned users. Figure 4 is an example email created shows that today's alert volume (blue dot) is higher than the upper bound (red dotted line), that is calculated from volume in the same day in last week.

Result

We implemented the dashboard at April in 2018 and started an alert optimization working group to review data. Table 2 shows descriptive statistics of drug duplicate alert generation during the proof-of-concept period (April to August 2018) and types of overridden reason entered (Table 3). Besides routine monitoring and discussion efforts the group has made, we introduced three cases of detecting abnormal pattern and optimizing alerts using the tools.

Table 2. Descriptive statistics

# of patient	183,448	# of alert dialog	637,071
# of patient visit	253,583	# of alert firing	2,068,790
# of provider alerted	14,621	# of overridden reason entered	213,226
# of facility / clinic	706	# of suppression	1,262,747
# of medication orders	10,916,693	# of alert dialog with behavioral change	41,123

Table 3. Overridden reason entered

Overridden reason type	#Record	Percentage
Prescriber Clinical Judgment	170,285	81%
Prescriber Consulted, OK Received	19,710	9%
Patient Already Tolerating	12,790	6%
Pharmacist Clinical Judgment	7,941	4%
Accept Previous Override Reason	22	0%
Total	210,748	100%

Case #1. Reducing nuisance alerts individually: With the combined information of mCDS end-user observation and effectiveness analysis from the dashboards, we added suppression for Dextrose 10%, 25%, 50% and 70% (3/22), and Humalog insulin (8/29). Figure 5 shows duplicate alert volume from the medications were dropped after the actions (red line).

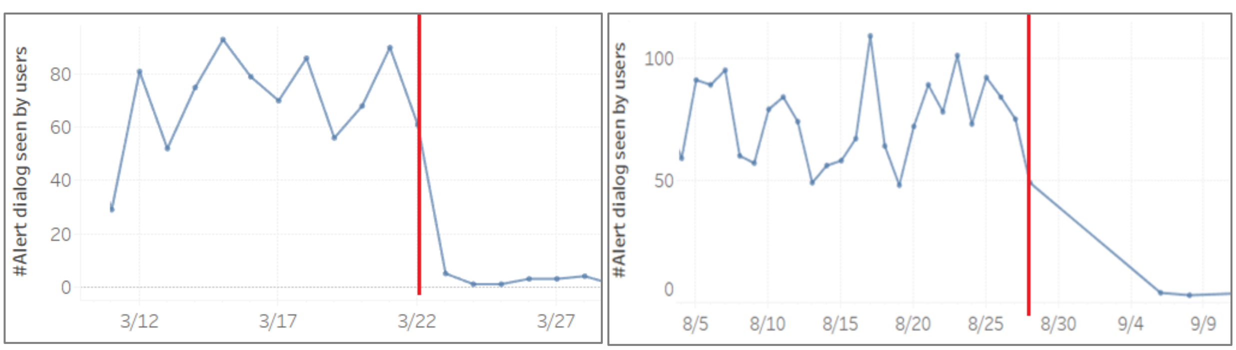

Figure 5. Reduction of duplicate alert: left) Dextrose 10%, 25%, 50% and 70%; right) Humalog (lispro) insulin

Case #2. Early detection of filtering failure for order set related duplicate alert: Our CPOE provides two ways for ordering medications, one as ordering individual medication items and the other as order sets that are a set of orders designed to treat patients for specific conditions or procedures. Order sets often cause duplicate alerts, since they contain a number of contents and orderable items and providers may place an order set without aware of some medications in it conflict already placed orders for a patient. Therefore we set up a baseline filtering so that providers do not receive duplicate alerts based on order checking against order sets. As of 4/27/2018, we detected abnormal duplicate alert peak in the dashboard and analyzed to figure out the filter for two order sets: anesthesia perioperative for adult and pediatric were broken. We recovered the filter at 5/2/2018 and the volume returned to normal (Figure 6). We received positive feedbacks from anesthesiologists, surgeons, nurses and pharmacists.

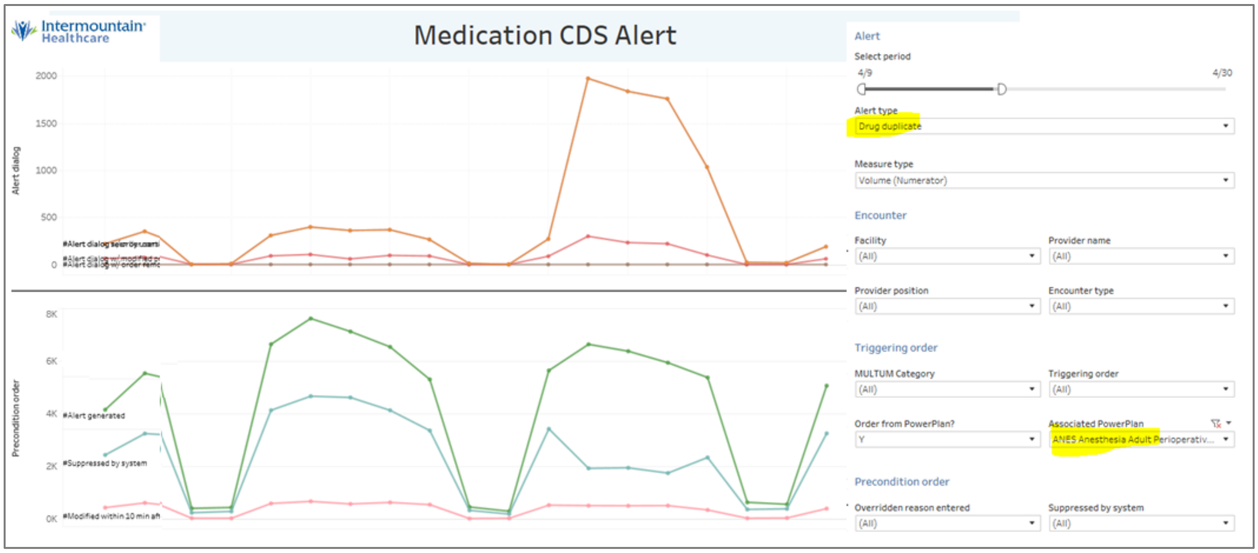

Figure 6. The dashboard provides a filter to drill down the volume of duplicate alerts affected by an order set and a suppression filter.

Case #3. Detecting broken queries in applications: We detected skyrocketed duplicate alerts seen by pharmacists at 4/19/2018. Through discussion and analysis internally and with other teams, we found that the reason was a custom query that filters out duplicate alerts when an order comes from CPOE to MedManger (pharmacist medication order verification) in the same ordering conversation. It was discovered that a new rule that was designed to run based on location was broken and interfering with the custom duplicate filtering CCL and causing it to not filter based on rule hierarchy. We discussed with a team creating the location rule and turned it off, the data showed a decrease in duplicate alerting (Figure 7).

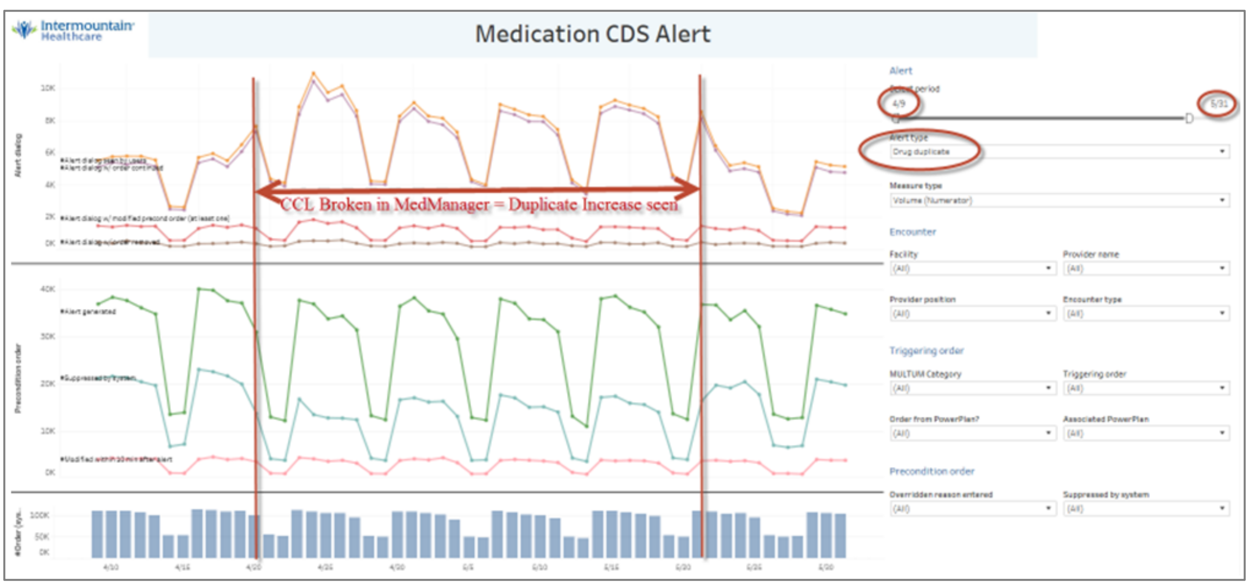

Figure 7. Increased volume of duplicate alerts during the period of broken query

Effective analysis: As a result of the efforts to optimize duplicate alerts for six months, the number of alert dialog presented has been decreased (Figure 8 left). Both raw and normalized volume were decreased with statistical significance (Table 4). We analyzed two effectiveness measures for the period, both behavioral change rate and overridden reason entered rate increased over time (Figure 8 right) with statistical significance (See Table 4).

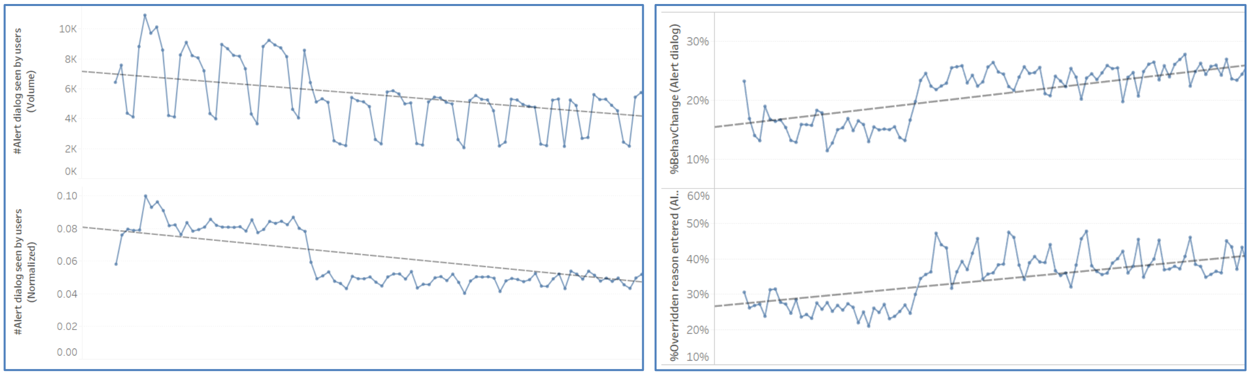

Figure 8. left) daily duplicate alert volume trend (top: volume, bottom: normalized volume); right) effectiveness metrics (top: % behavioral change, bottom: % overridden reason entered)

Table 4. Regression analysis: alert reduction over time; behavioral change / overridden rate over time

Measure	Line		Coefficients				
	p-value	Degree of freedom	Term	Value	Std. Error	t-value	p-value
#Alert dialog seen by users (Volume)	< 0.0001	111	Time	-31.4658	5.34406	-5.88799	< 0.0001
			intercept	1.36657e+06	231211	5.91047	< 0.0001
#Alert dialog seen by users (Normalized)	< 0.0001	111	Time	-0.0003536	3.058e-05	-11.5618	< 0.0001
			intercept	15.356	1.32307	11.6063	< 0.0001
%Behavioral change	< 0.0001	111	Time (day)	0.001024	8.437e-05	12.1376	< 0.0001
			Intercept	-44.0869	3.65014	-12.0781	< 0.0001
%Overridden reason entered	< 0.0001	111	Time (day)	0.001396	0.000149	9.36702	< 0.0001
			Intercept	-60.0532	6.4486	-9.31259	< 0.0001

Discussion and Conclusion

The results from proof-of-concept development and early usage analysis demonstrated that the proposed approach was successful to reduce drug duplicate alerts while maintaining their effectiveness. Although this is a case study with empirical data analysis, we will strive to generalize the proposed approach across other mCDS alert types: drug-drug interaction, allergy, dose checking, etc. In addition, we will develop detailed effectiveness metrics to more accurately measure how alerts affects to provider's behaviors and clinical processes.

This study has several limitations. Although we comprehensively evaluated effectiveness of alerts during the proof-of-concept period, it wasn't clearly investigated for how much individual actions affected to alert effectiveness. Beside our optimization efforts, there have been a number of administrative modifications done in the mCDS system, such as new rule definitions, drugs items, drug categories, and order sets. Because of this it was challenging to segregate alert reduction only affected by our optimization efforts.

In this study, we did not include clinical context of mCDS alerts into the analysis, such as patient encounter types, clinical condition, facilities, and provider positions. We found out that about half of duplicate alerts were seen by pharmacy and the rest by physicians. Since nuisance duplicate alerts used to occur between ordering providers and referred pharmacists, the interactive visual analytics approach will be useful to understand such patterns in the clinical processes.

Acknowledgement

The authors do not have any conflict of interest with the Cerner Corporation.

References

1. Wright A, Sittig DF. A four-phase model of the evolution of clinical decision support architectures, International Journal of Medical Informatics. 2008; 77(10): 641-649.
2. Kaushal R, Shojania KG, Bates DW. Effects of computerized physician order entry and clinical decision support systems on medication safety: a systematic review. Archive of Internal Medicine. 2003; 23;163(12):1409-16.

3. Schedlbauer A, Prasad V, Mulvaney C, Phansalkar S, Stanton W, Bates DW, Avery AJ. What evidence supports the use of computerized alerts and prompts to improve clinicians' prescribing behavior? Journal of the American Medical Informatics Association : JAMIA. 2009;16(4):531-538.
4. Page N, Baysari MT, Westbrook JI. A systematic review of the effectiveness of interruptive medication prescribing alerts in hospital CPOE systems to change prescriber behavior and improve patient safety. International Journal of Medical Informatics. 2017;105:22-30.
5. Nuckols TK, Smith-Spangler C, Morton SC, et al. The effectiveness of computerized order entry at reducing preventable adverse drug events and medication errors in hospital settings: a systematic review and meta-analysis. Systematic Reviews. 2014;3:56.
6. Van der Sijs H, Aarts J, van Gelder T, Berg M. Vulto A. Turning off frequently overridden drug alerts: Limited opportunities for doing it safely. Journal of the American Medical Informatics Association : JAMIA. 2008;15(4):439-448.
7. Nanji KC, Slight SP, Seger DL, et al. Overrides of medication-related clinical decision support alerts in outpatients. Journal of the American Medical Informatics Association : JAMIA, 2014;21(3):487-491.
8. Nanji KC, Seger DL, Slight SP, Amato MG, Beeler PE, Her QL, Dalleur O, Eguale T, Wong A, Silvers ER, Swerdloff M, Hussain ST, Maniam N, Fiskio JM, Dykes PC, Bates DW. Medication-related clinical decision support alert overrides in inpatients, Journal of the American Medical Informatics Association : JAMIA, 2018;25(5): 476–481.
9. Simpao AF, Ahumada LM, Desai BR, Bonafide CP, Gálvez JA, Rehman MA, Jawad AF, Palma KL, Shelov ED. Optimization of drug-drug interaction alert rules in a pediatric hospital's electronic health record system using a visual analytics dashboard. Journal of the American Medical Informatics Association : JAMIA. 2015;22(2):361-9.
10. Kwint HF, Faber A, Gussekloo J, Bouvy ML. Effects of medication review on drug-related problems in patients using automated drug-dispensing systems: a pragmatic randomized controlled study. Drugs Aging. 2011 Apr 1;28(4):305-14.
11. Heringa M, Floor-Schreudering A, Tromp PC, de Smet PA, Bouvy ML. Nature and frequency of drug therapy alerts generated by clinical decision support in community pharmacy. Pharmacoepidemiol Drug Saf. 2016 Jan;25(1):82-9.
12. McDaniel RB, Burlison JD, Baker DK, et al. Alert dwell time: introduction of a measure to evaluate interruptive clinical decision support alerts. Journal of the American Medical Informatics Association : JAMIA. 2016;23(e1):e138-e141.
13. Schreiber R, Gregoire JA, Shaha JE, Shaha SH. Think time: A novel approach to analysis of clinicians' behavior after reduction of drug-drug interaction alerts. International Journal of Medical Informatics, 2017;97:59-67.
14. Evans RS, Lloyd JF, Pierce LA. Clinical use of an enterprise data warehouse. AMIA Annual Symposium Proceedings, 2012;2012:189-198.